

How to Solve Problems

Chapter 1



STEP 1: Write down the answer [4]. In other words, come up with a reasonably close solution. This is frequently all the information you need.

For example, if it is 250 miles from New York to Boston, how long will it take to drive? You would immediately estimate that it should take about four or five hours, based on an average speed of 50–60 mph. This is enough information to decide whether or not you will drive to Boston for the weekend. If you do decide to drive, you will look at maps or the Internet and figure out the exact route and the exact expected driving time.

Similarly, before you go into a store, you usually know how much you are willing to spend. You might think it is reasonable to spend about \$100 on an X-Game2. If you see it for \$30, you will automatically buy it. If it sells for \$300, you will automatically not buy it. Only if the price is around \$100 will you have to think about whether to buy it.

We will apply the same reasoning here. We'll try to estimate the answer to within a factor of ten. Why a factor of ten? Because that is good enough to make most decisions.

Once you have estimated the answer to a problem, the answer will fall into one of the three “Goldilocks” categories:

1. too big
2. too small
3. just right

If the answer is too big or too small, then you know what to do (e.g., buy the item, don't drive to Boston). Only if the answer is just right will you need to put more work into solving the problem and refining the answer. (But that's beyond the scope of this book. We just aim to help you estimate the answer to within a factor of ten.)

If all problems were as simple as that, you wouldn't need this book. Many problems are too complicated for you to come up with an immediate correct answer. These problems will need to be broken down into smaller and smaller pieces. Eventually, the pieces will be small enough and simple enough that you can estimate an answer for each one. And so we come to

STEP 2: If you can't estimate the answer, break the problem into smaller pieces and estimate the answer for each one. You only need to estimate each answer to within a factor of ten. How hard can that be?

It is often easier to establish lower and upper bounds for a quantity than to estimate it directly. If we are trying to estimate, for example, how many circus clowns can fit into a Volkswagen Beetle, we know the answer must be more than one and less than 100. We could average the upper and lower bounds and use 50 for our estimate. This is not the best choice because it is a factor of 50 greater than our lower bound and only a factor of two lower than our upper bound.

Since we want our estimate to be the same factor away from our upper and lower bounds, we will use the geometric mean. To take the *approximate* geometric mean of any two numbers, just average their coefficients and average their exponents.* In the clown case, the geometric mean of one (10^0)[†] and 100 (10^2) is 10 (10^1) because one is the average of the exponents zero and two. Similarly, the geometric mean of 2×10^{15}

* We use coefficients and exponents to describe numbers in scientific notation. The exponent is the power of ten and the coefficient is the number (between 1 and 9.99) that multiplies the power of ten. If you are not familiar with this notation, please quickly check the section on scientific notation ("Dealing with Large Numbers") and then come right back. We'll wait for you here.

† Any number raised to the 0th power is 1.

and 6×10^3 is about 4×10^9 (because $4 = \frac{2+6}{2}$ and $9 = \frac{15+3}{2}$).^{*} If the sum of the exponents is odd, it is a little more complicated. Then you should decrease the exponent sum by one so it is even, and multiply the final answer by three. Therefore, the geometric mean of one and 10^3 is $3 \times 10^1 = 30$.

EXAMPLE 1: MongaMillions Lottery Ticket Stack

Here's a relatively straightforward example: Your chance of winning the MongaMillions lottery is one in 100 million.[†] If you stacked up all the possible different lottery tickets, how tall would this stack be? Which distance is this closest to: a tall building (100 m or 300 ft), a small mountain (1000 m), Mt Everest (10,000 m), the height of the atmosphere (10^5 m), the distance from New York to Chicago (10^6 m), the diameter of the Earth (10^7 m), or the distance to the moon (4×10^8 m)? Imagine trying to pick the single winning ticket from a stack this high.

Solution: To solve this problem, we need two pieces of information: the number of possible tickets and the thickness of each ticket. Because your chance of winning is one in 100 million, this means that there are 100 million (10^8) possible different tickets.[‡] We can't reliably estimate really thin items like a single lottery ticket (is it 1/16 in. or 1/64 in.? is it 1 mm or 0.1 mm?) so let's try to get the thickness of a pack of tickets.

^{*}To be more precise (which this book rarely is), the geometric mean of two numbers, b and c , is $a = \sqrt{bc}$. Our approximate rule is exact for the exponents and close enough for this book for the coefficients.

[†]Lottery billboards frequently have the odds of winning in very small print at the bottom.

[‡]100 million = 100,000,000 or 1 followed by eight (count them!) zeros. This can be written in scientific notation as 1×10^8 .

Let's think about packs of paper in general. One ream of copier or printer paper (500 sheets) is about 1.5 to 2 in. (or about 5 cm since 1 in. = 2.5 cm) but paper is thinner than lottery tickets. A pack of 52 playing cards is also about 1 cm. That's probably closer. This means that the thickness of one ticket is

$$\begin{aligned}t &= \frac{1 \text{ cm}}{52 \text{ tickets}} = 0.02 \frac{\text{cm}}{\text{ticket}} \times \frac{1 \text{ m}}{10^2 \text{ cm}} \\ &= 2 \times 10^{-4} \frac{\text{m}}{\text{ticket}}\end{aligned}$$

Therefore, the thickness of 10^8 tickets is

$$T = 2 \times 10^{-4} \frac{\text{m}}{\text{ticket}} \times 10^8 \text{ tickets} = 2 \times 10^4 \text{ m}$$

2×10^4 m is 20 kilometers or 20 km (which is about 15 miles since 1 mi = 1.6 km).

If stacked horizontally, it would take you four or five hours to walk that far.

If stacked vertically, it would be twice as high as Mt Everest (30,000 ft or 10 km) and twice as high as jumbo jets fly.

Now perhaps you used the thickness of regular paper so your stack is a few times shorter. Perhaps you used 1 mm per ticket so your stack is a few times taller. Does it really matter whether the stack is 10 km or 50 km? Either way, your chance of pulling the single winning ticket from that stack is pretty darn small.

EXAMPLE 2: Flighty Americans

These problems are great fun because, first, we are not looking for an exact answer, and second, there are many different ways of estimating the answer. Here is a slightly harder question with multiple solutions.

How many airplane flights do Americans take in one year?

We can estimate this from the top down or from the bottom up. We can start with the number of airports or with the number of Americans.

Solution 1: Start with the number of Americans and estimate how many plane flights each of us take per year. There are 3×10^8 Americans.* Most of us probably travel once a year (i.e., two flights) on vacation or business and a small fraction of us (say 10%) travel much more than that. This means that the number of flights per person per year is between two and four (so we'll use three). Therefore, the total number of flights per year is

$$\begin{aligned} N &= 3 \times 10^8 \text{ people} \times 3 \text{ flights/person-year} \\ &= 9 \times 10^8 \text{ passengers/year} \end{aligned}$$

Solution 2: Start with the number of airports and then estimate the flights per airport and the passengers per flight. There are several reasonable size airports in a medium-sized state (e.g., Virginia has Dulles, Reagan-National, Norfolk, Richmond, and Charlottesville; and Massachusetts has Boston and Springfield). If each of the fifty states has three airports then there are 150 airports in the US. Each airport can handle at most one flight every two minutes, which is 30 flights per hour or 500 flights per 16-hour day. Most airports will have many fewer flights than the maximum. Each airplane can hold between 50 and 250 passengers. This means

*This is one of those numbers you need to know to do many estimation questions. Go ahead and write it on your palm so you'll be prepared for the test at the end of the book.

that we have about

$$N = 150 \text{ airports} \times \frac{100 \text{ flights}}{\text{airport-day}} \times \frac{100 \text{ passengers}}{\text{flight}} \\ \times \frac{365 \text{ days}}{\text{year}} = 5 \times 10^8 \text{ passengers per year}$$

Wow! Both methods agree within a factor of two.

The actual number of US domestic airline passengers in 2005 was 6.6×10^8 , which is close enough to both answers.

EXAMPLE 3: Piano Tuners in Los Angeles

Now let's work out a harder problem.

How many piano tuners are there in Los Angeles (or New York or Virginia Beach or your own city)? This is the classic example originated by Enrico Fermi [5] and used at the beginning of many physics courses because it requires employing the methods and reasoning used to attack these problems but does not need any physics concepts.

Solution: This is a sufficiently complicated problem that we cannot just estimate the answer. To solve this, we need to break down the problem. We need to estimate (1) how many pianos there are in Los Angeles and (2) how many pianos each tuner can care for. To estimate the number of pianos, we need (1) the population of the city, (2) the proportion of people that own a piano, and (3) the number of schools, churches, etc. that also have pianos. To estimate the number of pianos each tuner can care for, we need to estimate (1) how often each piano is tuned, (2) how much time it takes to tune a piano, and (3) how much time a piano tuner spends tuning pianos.

This means that we need to estimate the following:

1. population of Los Angeles
2. proportion of pianos per person
3. how often each piano is tuned per year
4. how much time it takes to tune each piano
5. how much time each piano tuner works per year

Let's take it from the top.

1. The population of Los Angeles must be much less than 10^8 (since the population of the US is 3×10^8). It must be much more than 10^6 (since that is the size of an ordinary big city). We'll estimate it at 10^7 .
2. Pianos will be owned by individuals, schools, and houses of worship. About 10% of the population plays a musical instrument (it's surely more than 1% and less than 100%). At most 10% of musicians play the piano and not all of them own a piano so the proportion that own a piano is probably 2–3% of the musicians. This would be 2×10^{-3} of the population. There is about one house of worship per thousand people and each of those will have a piano. There is about one school per 500 students (or about 1 per 1000 population) and each of those will have a piano. This gives us about 4 or 5×10^{-3} pianos per person. Thus, the number of pianos will be about $10^7 \times 4 \times 10^{-3} = 4 \times 10^4$.
3. Pianos will be tuned less than once per month and more than once per decade. We'll estimate once per year.
4. It must take much more than 30 minutes and less than one day to tune a piano (assuming that it is not too badly out of tune). We'll estimate

2 hours. Another way to look at it is that there are 88 keys. At 1 minute per key, it will take 1.5 hours. At 2 minutes per key, it will take 3 hours.

5. A full-time worker works 8 hours per day, 5 days per week, and 50 weeks per year which gives $8 \times 5 \times 50 = 2000$ hours. In 2000 hours she can tune about 1000 pianos (wow!).

This means that the 4×10^4 pianos need 40 piano tuners.

How close are we? Well, the Yellow Pages for our city of 10^6 inhabitants (ten times fewer than LA) has 16 entries under the heading of “Pianos—Tuning, Repairing & Refinishing.” There are probably only one or two tuners per entry and they probably do not spend full time tuning. This means that our estimate is probably too low by a factor of five. However, that is a LOT closer than we could get by just guessing.

Remember that we are only trying to estimate the answer within a factor of ten.